

FASTER — A TOOL FOR DSN FORECASTING AND SCHEDULING

David Werntz, Steven Loyola, and Silvino Zencjas
Jet Propulsion Laboratory
Pasadena, California

Abstract

FASTER (Forecasting And Scheduling Tool for Earth-based Resources) is a suite of tools designed for forecasting and scheduling of JPL's Deep Space Network (DSN). The DSN is a set of antennas and other associated resources that must be scheduled for satellite communications, astronomy, maintenance, and testing. FASTER is a MS-Windows based program that replaces two existing programs (RALPH and PC4CAST). FASTER was designed to be more flexible, maintainable, and user friendly. FASTER makes heavy use of commercial software to allow for customization by users. FASTER implements scheduling as a two pass process: the first pass calculates a predictive profile of resource utilization; the second pass uses this information to calculate a cost function used in a dynamic programming optimization step. This information allows the scheduler to "look ahead" at activities that are not as yet scheduled. FASTER has succeeded in allowing wider access to data and tools, reducing the amount of effort expended and increasing the quality of analysis.

Background and History

FASTER, the Forecasting And Scheduling Tool for Earth-based Resources is a suite of software tools developed at JPL (Jet Propulsion Laboratory) to aid in the process of allocating DSN (Deep Space Network) 70 and 34 meter antennas and equipment to track deep space satellites (e.g., Galileo and Voyager) and ground based astronomy (e.g., Goldstone Solar System Radar). In addition, preventive maintenance, testing, and upgrades must be scheduled,

The 70 and 34 meter networks of the DSN presently include a total of 9 antennas (with plans for 3 to 9 more antennas in the next ten years). In addition, there are multiple transmitters, receivers and other pieces of equipment that must be coordinated with the antennas. These nine antennas are located at three sites around the world: Goldstone, California; Canberra, Australia; and Madrid, Spain (chosen to allow continuous coverage of deep-space objects near the ecliptic). Each site has three different type of antennas — a 70 meter (70 M), a 34 meter standard (34 S), and a 34 meter high efficiency (34H).

Normal activities that utilize the antennas (i.e. not maintenance type activities) generally have three parts:

- *Pre-calibration* — this time is dedicated to configuring the antenna, testing the configuration, and calibrating equipment. This may occur while the object to be tracked is not in view. Typical pre-calibration times run between 45 minutes and one and one half hours. The time required is determined by the type of tracking to be performed and the equipment required.
- *Track* — this is the time during which actual transmission and/or reception occurs. The object to be tracked must be in view during this time. Typical times range from two to ten hours. The amount of time required is determined by data rates, data quantity, desired risk, and many other factors.
- *Post-Calibration* — this time is dedicated to de-configuring the antenna and its associated equipment. Typical times range from 15 minutes to 45 minutes and are determined by the configuration. This may occur while the object is not in view

Under normal circumstances these three parts occur consecutively with no intervening time. Taken together, these three parts comprise a *pass*; however, we will often use the generic term *activity* to mean a pass or other use of an antenna and equipment (e.g., maintenance),

The previous paragraph referred to an object as being *in view*. This implies that an object is "visible" from the antenna. Each time during which an object is in view is known as a view *period*. A view period is said to have a rise (the start of the view period) and a set (the end of the view period). The calculation of view periods also involves information about surrounding terrain, antenna geometry, and signal to noise limitations. These calculations are performed by a program external to FASTER and stored in a database. For the majority of the deep space objects, this all translates to having one, approximately 10 hour view period per site per day that varies seasonally (in time of day and duration).

Schedules are segmented into weekly pieces. Initial schedules, with conflicts in them, are generated six months to two years in advance of real time. This allows time for negotiation of conflicts and for other more detailed planning functions to occur (e.g. detailed spacecraft sequence generation). In general, schedules are

resolved until they are conflict free, two months in advance of real time.

Currently, approximately 200 to 300 activities are scheduled weekly for these antennas and the numbers continue to grow. On average, only 50-60% of requirements are met and any improvement in scheduling efficiency can result in significant additional science return.

While the goal in scheduling is to generate schedules that utilize the resources as efficiently as possible, the goal in forecasting is to help set the stage so that scheduling is easier and more efficient. Forecasting can be thought of as a "what if" capability. Someone poses a change and would like to know the impact of that change. While it is true that one could generate schedules and evaluate schedules, it is not necessarily feasible. Often, the level of detail necessary to generate a schedule does not exist (e.g., in the case of a spacecraft that is just in the design phase). Another problem is that the change should be evaluated over a fairly long period of time (years). Even if schedules could be created at the rate of one per minute, a four year study would take over four hours to generate. Finally, it can be difficult to generalize from a schedule; it is very easy to infer data from a detailed schedule that is really only an artifact of that particular schedule solution.

There are a variety of questions that are posed as "what-if" studies; the processing of these studies has become a major consumer of effort. Some examples include:

- What would be the impact of cutting pre-calibration time in half?
- If another antenna were to be added to the network, at which of the three sites should it be added to do the most good?
- If project A were funded, do we have enough resources to support it and all other missions that are already funded? If not, what might get cut?

In the area of long term forecasting, lead times for both resource augmentation and project (spacecraft) development are measured in years and costs are in 100's of millions of dollars; therefore, it is of utmost importance to properly answer these questions. Currently, forecasting is performed as far out as 20 years and as close to real time as two years.

Since 1986, the RALPH (Resource Allocation and Planning Helper) system has been operationally used to assist in the scheduling process. RALPH runs on VAX hardware under VMS and has primarily a character-based

interface for use on dumb terminals. In 1989, a prototype, called PC4CAST, was developed to assist in forecasting tasks. PC4CAST ran on PC compatible hardware under DOS and used Quattro Pro (a spreadsheet) for tabular data entry and for results graphing. The majority of calculations were done by an external program, reading and writing Quattro data files. RALPH and PC4CAST had no linkage and because of the different platforms, some data had to be duplicated.

Algorithms

There are two algorithms of interest in FASTER, generically known as the first and second pass algorithms. The scheduling tool uses both algorithms, while the forecasting tool uses only the first pass. In the first pass, FASTER creates a set of *expected usage profiles*, which represent a statistical analysis of resource demand. The second pass uses this information to derive a cost function which drives a dynamic programming algorithm for group activity scheduling. Group activity scheduling implies that activities are not scheduled individually but rather that a set of relatively homogeneous activities are scheduled as a whole. A group of activities is often referred to as one request or project requirement. A typical requirement might be that Pioneer 10 wants seven, four to eight hours tracks on any 70M antenna and the tracks must be separated by at least 10 hours and by no more than 36 hours.

The forecasting tool uses the *expected usage profiles* from the first pass algorithm to generate statistics that are both consistent with scheduling and designed to give insight into resource and requirement problems. Many statistics can be derived, including: expected lost time and number of resources required to meet performance requirements.

First Pass (Forecast)

In forecasting, many details are presently ignored that would be considered crucial to scheduling; however, the goal is to gain as much insight into scheduling problems without resorting to the level of detailed data input required for scheduling or the attendant run-time. For example, the definition of a request for forecasting is considerably simpler than that for scheduling. A forecasting request is defined by:

- View period object (e.g., spacecraft, planet, shifts)
- Usable antenna resources
- Pre- and post-calibration times
- Minimum track duration
- Average Track duration
- Number of tracks per week

The end result of the first pass is a set of expected usage profiles. One profile is generated for each resource. Each profile represents the expected usage of that resource as a function of time. In addition, expected usage values are subdivided by activity group as well. Expected usage profiles are constructed by looping through each group of activities that is to be scheduled.

The calculation of expected usage profiles is a three stage process that incorporates requirements, resource capacity, and view period information. The steps are:

- 1) Calculate an *expected usage value* for each requirement. This value represents the ratio of time that must be scheduled for a requirement to the time that is available for the requirement to be scheduled.
- 2) The expected usage value is used with the view periods and the pre- and post-calibration times to generate individual expected usage profiles. Each of these profiles represents the demand from one requirement for each point in time for each antenna.
- 3) Individual expected usage profiles for each antenna are summed, resulting in one expected usage profile for each antenna.

Each of these steps is described in greater detail below.

Calculation of Expected Usage Value

In calculating the expected usage value for a requirement, first the total amount of requested time is calculated. This is calculated based upon the requirement's average duration, pre-, and post-calibration times, and the number of tracks requested. Calibration times are included because they represent demand on the resources just as the actual track.

The next step is to find all usable view periods for the specified set of resources that are long enough to support the minimum requested duration. Then, the pre- and post-calibration times are appended to them, resulting in what is called *request slots*. Request slot time is defined as the sum of all durations of these request slots. An expected usage is then calculated and represents the total requested time divided by total request slot time. If the expected usage is greater than one, this means that the requirement cannot be supported by the resource(s) specified. This is caused by physical constraints or resource downtime.

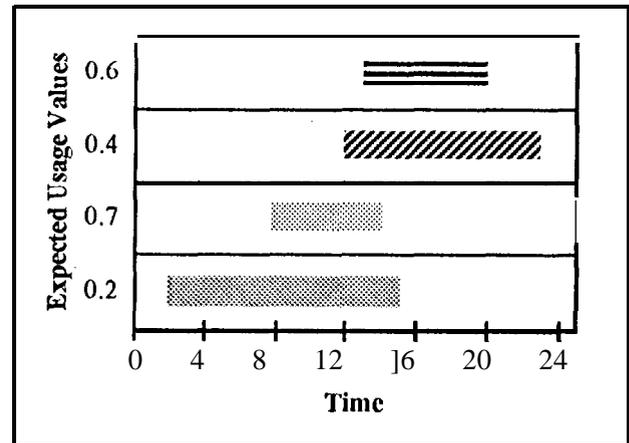


Figure 1- Example Request Slots and Usage Values

Calculation of Individual Profiles

The next step is to generate individual expected usage profiles. In FASTER, these profiles are represented by step functions. The profile for a single request will be a two level step function having the request's expected usage during all usable request slots and the value zero at all other times.

Calculation of Antenna Profiles

The last step is to sum all individual request profiles for each antenna. This results in a complete picture of the expected usage of each antenna at all times. By using this information to drive the scheduling phase, the scheduler is able to determine what areas should be avoided even when nothing has actually been scheduled in that area. A summed profile dissection is shown in figure 2 (based upon data from figure 1). An example of a complete one week antenna profile is shown in figure 3.

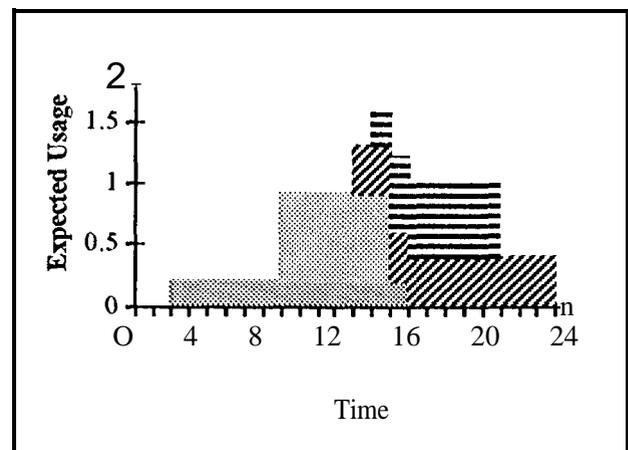


Figure 2- Example Sumation of Individual Profiles

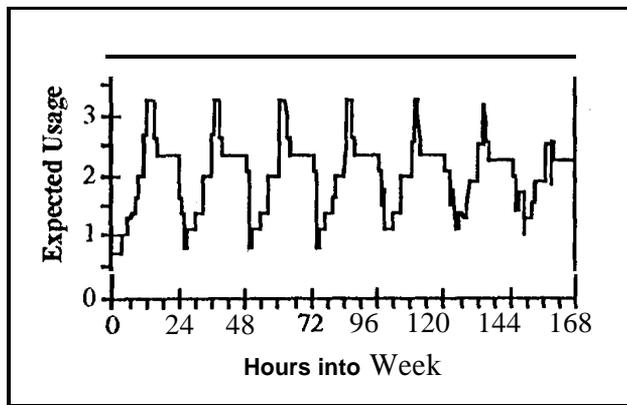


Figure 3- Example Antenna Expected Usage Profile

Second Pass (Schedule)

FASTER performs a series of **optimizations** steps, **one** for each requirement that is scheduled. Decomposing a **problem** into multiple optimization steps does not necessarily result in a solution that is anywhere near **optimal**. To mitigate the **effects** of the **suboptimization**, the expected usage profiles are used to derive the cost function for each optimization and thus **allow** the **scheduler** to “look ahead” at requests that have yet to be scheduled.

As mentioned previously, the description of a **request** is more detailed for scheduling. In particular, separation constraints (minimum, maximum, and preferred) are added to force a desired **distribution** of tracks. Separation constraints determine the time between the **end** of one track in a request and the start of the next track in the request.

The **first** step in scheduling a request is to remove its **contribution** from all expected usage profiles. Next a set of **break points** are calculated, indicating all points at which it would be reasonable to begin or end a track. They are called break points **because** they correlate to breaks in the expected usage profiles. Break points are distinguished as to whether they are start or end break points. This approach transforms the problem from finding the best place to schedule n tracks to finding a minimum cost “path” through the correct number of break points such that all duration and separation constraints are met.

This **problem** is similar to that of the classical traveling salesman problem where a shortest path must be found through a set of cities. ¹ Instead of finding a path through cities, find a path through alternating start and **end** break points.

When traveling from a start break point to an end break points, the following **must** hold:

- the **end** point being considered **must** have come from the same view period as the start point
- the **end** point must **be** at least the minimum duration after the start point
- the **end** point must **be** less than the maximum duration after the **start** point

The cost for a valid start to end path (a track) is the combination of a duration preference value and an **expected** resource cost. The expected resource cost is **calculated** by integrating the area under the expected usage profile for the chosen antenna during the interval **start-pre** to **end+post**.

When traveling from an end break point to a start break point, the following must hold

- the start must **be** at least the minimum separation after the end
- the start must **be** less than the maximum duration after the end

The cost for a valid end to start path (gap) is calculated simply from any preferences on separation (since no **resources** are consumed during this time).

In addition, several other properties of the **problem** can **be** used to reduce the total number of paths that must **be** evaluated. For example, given that all of the tracks must fit within a certain period of time (an overall window), upper and lower bounds can be calculated for the start and **end** time of each track. Another example is to calculate a simple first-cut path and use its cost to **prune** paths as they are calculated.

Once the optimal path has been calculated, the tracks of that request are actually scheduled and the expected usage profiles are updated to reflect the new information. Since the **antenna** and start and end **times** for the tracks are now **known**, a contribution of 1.0 is added in for each track. In this way, the scheduling pass can **be** viewed as transforming the expected usage profiles from probabilistic to deterministic.

Limitations

While this approach to scheduling and forecasting has been successful in our particular **problem**, there are some limitations to the current approach that might limit its usefulness in other cases. These include:

- The need for “view period like” restrictions — If there are few or no limitations on where activities

may occur, profiles will be of little use (a flat profile is uninteresting).

- The need for durations and numbers of activities to be known prior to scheduling (within ranges) — If the total amount of time that might be scheduled for a request can vary over an order of magnitude range, then calculation of expected usage values may be difficult or extremely inaccurate,
- The need for an oversubscribed problem with requirement interaction — if requirements do not overlap sufficiently, enough break points may not be generated and feasible schedules may not be found.

Improvements

As with all solutions, there are steps that could be taken to improve upon the approach. These include:

- Better search algorithms — little attention has been paid to reducing the number of paths that are evaluated.
- Better modeling of probability distribution — currently profiles are modeled as step functions; this is not wholly accurate. A true representation would involve piece-wise linear functions. However, it is not clear that the additional overhead would really improve schedule quality.
- Consider spacing constraints in the generation of expected usage profiles — the current profile generation scheme implies a more or less even distribution of activities. Depending upon view periods and spacing interaction, this may not be a good assumption.
- Implement an incremental update feature — the current system does not support incremental requirement modifications/additions.

Implementation

The FASTER system runs on a 386/486 class machine with at least 8MB of RAM running Microsoft Windows version 3.1.2. FASTER requires connection to a server running Novell Netware. This server is used both for program sharing as well as being the database server in a client/server architecture. FASTER is implemented in C++ and based upon a set of classes named RASCL (Resource Allocation and Scheduling Class Library). RASCL was codeveloped by JPL and CTA (a Langley contractor). RASCL implements many of the common data structures and primitives used in scheduling and forecasting, including:

- integer time representation
- intervals
- timelines
- profiles

The forecasting portion of FASTER uses Microsoft Excel (a spreadsheet) as a tabular input sheet and for graphing statistical information. Microsoft Access is being used as a client front end to the database server for database maintenance and report generation. The use of these COTS (commercial off the shelf) packages has increased programmer productivity, allowed users to customize reports and graphs, and reduced the maintenance burden.

In terms of performance, the forecasting tool is able to generate an average one year forecast on a 33MHz 486 in under five minutes. This includes the time to retrieve approximately 30,000 view period records over the network, calculate expected usage profiles, calculate desired statistics from the profiles, and write a file with the desired statistical information for loading back into Excel.

Issues

For those who would be interested in trying some of the approaches we have described, in an operational environment, these issues should be addressed.

While using COTS packages has great advantages, it also can imply dependence upon that package. In particular: When a new version comes out, do you upgrade? Are there new capabilities that really would be useful? Is it really backward compatible? Does the compatibility require some sort of "conversion". Can everyone be upgraded at once? If you don't upgrade, can you still get technical support? As you would guess, the only way to really answer some of these questions is to try it. Install a new version in an off-line simulation of operations and test, test, test.

As our user base has grown and the development team has shrunk, training and support have become prime issues. When considering a large system, the normal questions must be asked. Who do you train? Do you charge for training? Who do you support? Do you charge for support? Since the program makes heavy use of a particular COTS package, do you have to be the support for the COTS package. In general, this has to be worked out with the customer and the funding source; just remember, customer support can be a differentiating factor when competing.

At JPL, as at other large companies, the demand for network services sometimes outstrips the supply.

FASTER is dependent upon the institutional network for users who are outside the local area of the server. Network problems can become a support headache that you may be near powerless to affect. Do some testing to determine performance and become involved in any planning committees that address these types of problems.

Acknowledgments

The authors wish to thank Eugene Burke of the Multimission Operations Systems Office for his continued support of FASTER development and J.R. Hall of the Office of Telecommunications and Data Acquisition for his continued support of forecasting development. U.S. Government sponsorship under contract NAS7-918 is acknowledged.

References

- ¹Cormen, Thomas H., Charles E. Leiserson and Ronald L. Rivest, *Introduction to Algorithms*, MIT Press, 1990.
- ²Wertz, David, Silvino Zendejas and Steven Loyola, *Forecasting And Scheduling Tool for Earth-based Resources (FASTER) User Guide, 1992*, JPL Internal Document D-10450.